Performance Analysis of Applications under CPU Power Constraints

Riki Takahashi
takahashi.riki.66a@st.kyoto-u.ac.jp
Kyoto University
Japan, Kyoto

Keiichiro Fukazawa
takahashi.riki.66a@st.kyoto-u.ac.jp
Kyoto University
Japan, Kyoto

1 INTRODUCTION

As the computational performance of supercomputers has improved, their power consumption has been escalating annually. Consequently, improving the energy efficiency of supercomputing systems has become one of the most critical challenges. In response to this issue, studies have been undertaken that consider job scheduling and the performance variations of processors, leading to optimizations at the middleware level. To further enhance power efficiency, it is essential not only to continue these middleware optimizations but also to build upon application-level optimizations.

To enhance power efficiency at the application level, it's crucial to optimize the performance of applications within the allocated power budget. While power modeling for supercomputers and performance estimation for applications have been undertaken [3], there is scarcely any analysis on application performance with power constraints.

For this problem, execution times under power constraints for various benchmark applications and simulation codes have been estimated [2]. In this study, the application performance under power constraints is modeled using the execution results at the CPU's minimum and maximum frequencies. Additionally, the performance during power constraints of both CPU and DRAM in a Magnetohydrodynamic (MHD) simulation code has been evaluated [1] . While these studies investigate application performance under power constraints, they focus on overall application performance and do not explore detailed performance such as Byte per Flops (B/F) values and power performance in specific sections of the application. However, actual applications are composed of various sections, each with distinct power characteristics, such as memory-bound or CPU-bound sections. Measuring the performance of applications in each section under power constraints is necessary for more efficient power control.

2 RESEARCH APPROACH

This study has done the following. Firstly, the research measured and evaluated the performance differences under CPU power constraints using the memory-bound benchmark application and the CPU-bound application. Secondly, it evaluated the relationship between B/F values and execution performance under CPU power constraints in practical scientific computing applications. Finally, the study conducted comprehensive measurements of execution performance in scientific computing applications, not only in general but also in specific parts like main computation and data writing, to understand the variations in power performance. Notably, this investigation focused on the power consumption of a single node.

3 EXPERIMENTAL ENVIRONMENT

This study was conducted using a single node from the Kyoto University Supercomputer System B. The node is configured with two sockets, each containing an Intel Xeon Platinum 8480+ (codenamed Sapphire Rapids) processor with 56 cores and DDR5 memory. The applications used in the experiments are simulation codes for scientific and technical computing, including Finite Difference Time Domain (FDTD), MHD, Particle-in-Cell (PIC), and several HPC challenge benchmarks. For power constraints and performance data acquisition, the RIC [2] library based on the Intel RAPL interface was utilized.

4 MEASUREMENT RESULTS

In our evaluation results of benchmark applications' performance under only CPU power constraints, we found that memory-bound applications tend to maintain their performance better than CPU-bound applications. In practical scientific computing applications, those with higher B/F values tend to have less performance reduction under CPU power constraints. Furthermore, by measuring execution characteristics under power constraints in different sections of scientific computing applications, we clarified the variations in power characteristics between these sections. These results suggest that finer power control within the application, based on sectional performance information, can further enhance power efficiency.

5 SUMMARY

This study addressed the power issues of supercomputers. The optimization of power control at the application level had not been extensively explored, but it became an increasingly important topic. In this research, we measured the performance of various applications under power constraints, and evaluated the relationship between the characteristics of these applications and their power performance, as well as the power characteristics within different sections of the applications. Measurement results indicate that power efficiency can be improved by implementing power control based on detailed performance such as B/F value of each application and different sections within the applications.

REFERENCES

[1] Keiichiro Fukazawa, Masatsugu Ueda, Yuichi Inadomi, Mutsumi Aoyagi, Takayuki Umeda, and Koji Inoue. 2018. Performance Analysis of CPU and DRAM Power Constrained Systems with Magnetohydrodynamic Simulation Code. (2018), 626–631. https://doi.org/10.1109/HPCC/SmartCity/DSS.2018.00113

[2] Yuichi Inatomi, Yuta Kakibuka, Taketsugu Ono, Kouji Inoue, et al. 2016. Performance Modeling in Power-Constrained Supercomputers [Translated from Japanese]. *Research Reports on High Performance Computing* 2016, 17 (2016), 1–9. .

[3] Kenneth O'brien, Ilia Pietri, Ravi Reddy, Alexey Lastovetsky, and Rizos Sakellariou. 2017. A survey of power and energy predictive models in HPC systems and applications. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 1–38.